


COMP 150

Writing a paper review

January 23, 2007



Business

- Get onto the mailing list
`comp150-ip1@eecs.tufts.edu`
- Start thinking about projects
 - Preferably in groups of 2 or 3
 - I'll provide some ideas
- Today:
 - Overview of paper reviewing
 - Look at papers



Tufts University Computer Science

1

Peer review

- This is how science works
 - Research papers submitted to conferences, journals, etc.
 - No central authority
 - Other researchers evaluate papers
 - Conference chair or editor sends out drafts to referees
 - Collect 3-6 reviews of a paper
 - Committee weighs the reviews, makes a decision
 - Reviews are anonymous
- One of the duties of a researcher



Tufts University Computer Science

2

Impact of reviewing

- Bad reviews break the process
 - Let bad work in – lowering the quality
 - Keep good work out
 - Delay dissemination of ideas
 - Discourage other researchers
 - Damage careers
- Writing good reviews requires practice
 - No formal definition of a "good" review
 - Typically, no training



Tufts University Computer Science

3

In this class

- Practice reviewing
 - Write a short review of each paper
 - Email it to me before class
 - Bring your issues and questions to class
- Note:
 - These are all published papers
 - But that doesn't mean they are perfect



Tufts University Computer Science

4

General guidelines

- Tone of a review
 - Don't hide behind anonymity
 - Write a review as if you are sitting down with the authors to discuss their work
- Be constructive
 - What are the good things about the paper?
- Be brutally honest
 - If something is bad, say it's bad
- Yes, it's a fine line



Tufts University Computer Science

5

Kinds of papers

- Evaluation differs depending on area
 - Theory paper
 - Checking mathematics
 - Soundness, correctness
 - Systems paper
 - Experimental methodology
 - Empirical results
- Much of the following applies to both



Core questions

- What is the problem being addressed?
- Is the goal significant?
- Does the work solve the problem?
- Is it novel?
- Is the methodology valid?
- Are the correct conclusions drawn from the results?
- Is the presentation satisfactory?



Motivation

- What is the problem?
 - *"This paper presents a new cache design in which the compiler directly controls which data items reside in cache"*
 - Is it a "real" problem?
 - Why do I care?
 - Who needs a solution to this problem?
- What are the assumptions?
 - *"CPU speeds will continue to grow faster than memory speeds."*
 - Technological assumptions – current and future
 - Economic factors
 - Overall trends in hardware, software, etc.



Solution

- Does it solve the problem? (duh)
 - Not always obvious – sometimes the goal is overstated
 - *"CPU speeds will continue to grow faster than memory speeds. We present a compiler-controlled cache..."*
- What are the contributions?
 - How is this work different and better than previous work?
 - (We'll come back to related work)
- Is it clearly explained?
 - System design
 - Algorithms and pseudo-code



Technical merit

- Is it a "good" solution?
 - What is the complexity of the solution?
 - Example: if we relax or modify the assumptions, is there a simpler solution?
 - Caveat: you can't object because you don't *like* the solution
 - *"We add two bits to each cache line..."*
- Is it implemented in the right place?
 - Classic tradeoff: hardware vs software
- What are the downsides?
 - Does it require complex engineering?
 - How easily will it coexist with other system components?



Evaluation

- Trickiest part of any research project
- Key question:
 - Do the experiments test and validate the hypothesis
- Specifics: the methodology
 - Is the methodology clearly described?
 - How is improvement measured?
 - What is the metric?
 - What is the baseline?
 - Details:
 - System platform: what are the relevant features?
 - Idiosyncratic features
 - What are the benchmark programs?



Related work

- Critical part of any paper
 - Places the work in context
 - Gives credit for previous contributions
 - Also: provides a citation chain – a history of the work
 - Two main kinds of related work
 - Work the paper builds upon
 - Alternative solutions to the same problem
 - What you're looking for:
 - The authors understand how the work is related
 - Brief description of how their work differs
- "Previous work has studied cache architectures [1, 2, 17, 18, 23, 29, 44, 49, 56]"*



Some pitfalls

- Radically new ideas
 - Hard to evaluate
- Moderately new ideas
 - Balance new idea against limited evaluation
- Old idea reformulated in a new environment
 - Example: Garbage collection in Java
 - Is the new work new enough?
- Everything sucks
 - Easy to hammer any paper



Real review

- This is an interesting idea and reasonably well-executed. However, the results are not especially compelling: performance improvements in very large heaps are easy to come by.
- Furthermore, I am puzzled by/suspicious of your choice of baseline collectors to measure. The Appel collector (variable nursery) typically outperforms the others; why didn't you implement your technique in this one, as it's available in Jikes RVM?
- I didn't find the measurements very helpful in understanding your technique. How does it affect the number of nursery vs. the number of mature space collections? Is that what determines performance? Or is it locality effects?
- There also needs to be some serious discussion of how this would be implemented in a JIT, rather than as an offline analysis.
- In the measurements section there is considerable confusion between "colocation" and "pre-nursing". You seem to use the terms interchangeably.
- Although I can't produce the reference, I could swear I've seen the colocation idea before, in the form of an allocator that takes an extra parameter which determines where to place the object.



For this class...

- One paragraph summary
 - Show that you understand the idea
- Strengths of the paper
 - Good idea? Strong implementation? Well-written?
 - What have the authors done right
- Weaknesses of the paper
 - Weak evaluation? Problem is irrelevant?
 - What can the authors do to improve the paper?



Today

- Look at papers to read...

